

Auto-Sizing the Transformer Network: Improving Speed, Efficiency, and Performance for Low-Resource Machine Translation

Kenton Murray Jeffery Kinnison Toan Q. Nguyen Walter Scheirer David Chiang

Department of Computer Science and Engineering

University of Notre Dame

{kmurray4, jkinniso, tnguye28, walter.scheirer, dchiang}@nd.edu

Abstract

Neural sequence-to-sequence models, particularly the Transformer, are the state of the art in machine translation. Yet these neural networks are very sensitive to architecture and hyperparameter settings. Optimizing these settings by grid or random search is computationally expensive because it requires many training runs. In this paper, we incorporate architecture search into a single training run through *auto-sizing*, which uses regularization to delete neurons in a network over the course of training. On very low-resource language pairs, we show that auto-sizing can improve BLEU scores by up to 3.9 points while removing one-third of the parameters from the model.

1 Introduction

Encoder-decoder based neural network models are the state-of-the-art in machine translation. However, these models are very dependent on selecting optimal hyperparameters and architectures. This problem is exacerbated in very low-resource data settings where the potential to overfit is high. Unfortunately, these searches are computationally expensive. For instance, Britz et al. (2017) used over 250,000 GPU hours to compare various recurrent neural network based encoders and decoders for machine translation. Strubell et al. (2019) demonstrated the neural architecture search for a large NLP model emits over four times the carbon dioxide relative to a car over its entire lifetime.

Unfortunately, optimal settings are highly dependent on both the model and the task, which means that this process must be repeated often. As a case in point, the Transformer architecture has become the best performing encoder-decoder model for machine translation (Vaswani et al., 2017), displacing RNN-based models (Bahdanau et al., 2015) along with much conventional wisdom about how to train such models. Vaswani

et al. ran experiments varying numerous hyperparameters of the Transformer, but only on high-resource datasets among linguistically similar languages. Popel and Bojar (2018) explored ways to train Transformer networks, but only on a high-resource dataset in one language pair. Less work has been devoted to finding best practices for smaller datasets and linguistically divergent language pairs.

In this paper, we apply *auto-sizing* (Murray and Chiang, 2015), which is a type of architecture search conducted during training, to the Transformer. We show that it is effective on very low-resource datasets and can reduce model size significantly, while being substantially faster than other architecture search methods. We make three main contributions.

1. We demonstrate the effectiveness of auto-sizing on the Transformer network by significantly reducing model size, even though the number of parameters in the Transformer is orders of magnitude larger than previous natural language processing applications of auto-sizing.
2. We demonstrate the effectiveness of auto-sizing on translation quality in very low-resource settings. On four out of five language pairs, we obtain improvements in BLEU over a recommended low-resource baseline architecture. Furthermore, we are able to do so an order of magnitude faster than random search.
3. We release GPU-enabled implementations of proximal operators used for auto-sizing. Previous authors (Boyd et al., 2010; Duchi et al., 2008) have given efficient algorithms, but they don't necessarily parallelize well on GPUs. Our variations are optimized for GPUs and are implemented as a general toolkit and are released as open-source software.¹

¹<https://github.com/KentonMurray/ProxGradPytorch>

2 Hyperparameter Search

While the parameters of a neural network are optimized by gradient-based training methods, hyperparameters are values that are typically fixed before training begins, such as layer sizes and learning rates, and can strongly influence the outcome of training. Hyperparameter optimization is a search over the possible choices of hyperparameters for a neural network, with the objective of minimizing some cost function (e.g., error, time to convergence, etc.). Hyperparameters may be selected using a variety of methods, most often manual tuning, grid search (Duan and Keerthi, 2005), or random search (Bergstra and Bengio, 2012). Other methods, such as Bayesian optimization (Bergstra et al., 2011; Snoek et al., 2012), genetic algorithms (Benardos and Vosniakos, 2007; Friedrichs and Igel, 2005; Vose et al., 2019), and hypergradient updates (Maclaurin et al., 2015), attempt to direct the selection process based on the objective function. All of these methods require training a large number of networks with different hyperparameter settings.

In this work, we focus on a type of hyperparameter optimization called auto-sizing introduced by Murray and Chiang (2015) which only requires training one network once. Auto-sizing focuses on driving groups of weights in a parameter tensor to zero through regularization. Murray and Chiang (2015) focused on the narrow case of two hidden layers in a feed-forward neural network with a rectified linear unit activation. In this work, we look at the broader case of all of the non-embedding parameter matrices in the encoder and decoder of the Transformer network.

3 GPU Optimized Proximal Gradient Descent

Murray and Chiang (2015) train a neural network while using a regularizer to prune units from the network, minimizing:

$$\mathcal{L} = - \sum_{f, e \text{ in data}} \log P(e | f; W) + \lambda R(\|W\|),$$

where W are the parameters of the model and R is a regularizer. For simplicity, assume that the parameters form a single matrix W of weights. Murray

Algorithm 1 Parallel ℓ_∞ proximal step

Require: Vector \mathbf{v} with n elements

Ensure: Decrease the largest absolute value in \mathbf{v} until the total decrease is $\eta\lambda$

- 1: $v_i \leftarrow |v_i|$
 - 2: sort \mathbf{v} in decreasing order
 - 3: $\delta_i \leftarrow v_i - v_{i+1}, \delta_n \leftarrow v_n$
 - 4: $c_i \leftarrow \sum_{i'=1}^i i' \delta_{i'}$ ▷ prefix sum
 - 5: $b_i = \frac{1}{i} (\text{clip}_{[c_{i-1}, c_i]}(\eta\lambda) - c_{i-1})$
 - 6: $p_i = \sum_{i'=i}^n b_{i'}$ ▷ suffix sum
 - 7: $\mathbf{v} \leftarrow \mathbf{v} - \mathbf{p}$
 - 8: restore order and signs of \mathbf{v}
-




Figure 1: Illustration of Algorithm 1. The shaded area, here with value $\eta\lambda = 2$, represents how much the ℓ_∞ proximal step will remove from a sorted vector.

and Chiang (2015) try two regularizers:

$$R(W) = \sum_i \left(\sum_j W_{ij}^2 \right)^{\frac{1}{2}} \quad (\ell_{2,1})$$

$$R(W) = \sum_i \max_j |W_{ij}| \quad (\ell_{\infty,1})$$

The optimization is done using proximal gradient descent (Parikh and Boyd, 2014), which alternates between stochastic gradient descent steps and proximal steps:

$$W \leftarrow W - \eta \nabla \log P(e | f; w)$$

$$W \leftarrow \arg \min_{W'} \left(\frac{1}{2\eta} \|W - W'\|^2 + R(W') \right)$$

To perform the proximal step for the $\ell_{\infty,1}$ norm, they rely on a quickselect-like algorithm that runs in $O(n)$ time (Duchi et al., 2008). However, this algorithm does not parallelize well. Instead, we use Algorithm 1, which is similar to that of Quattoni et al. (2009), on each row of W .

The algorithm starts by taking the absolute value of each entry and sorting the entries in decreasing order. Figure 1a shows a histogram of sorted absolute values of an example \mathbf{v} . Intuitively, the goal of the algorithm is to cut a piece off the top with area $\eta\lambda$ (in the figure, shaded gray).

We can also imagine the same shape as a stack of horizontal layers (Figure 1b), each i wide and δ_i high, with area $i\delta_i$; then c_i is the cumulative area of the top i layers. This view makes it easier to compute where the cutoff should be. Let k be the index such that $\eta\lambda$ lies between c_{k-1} and c_k . Then $b_i = \delta_i$ for $i < k$; $b_k = \frac{1}{k}(\eta\lambda - c_{k-1})$; and $b_i = 0$ for $i > k$. In other words, b_i is how much height of the i th layer should be cut off.

Finally, returning to Figure 1b, p_i is the amount by which v_i should be decreased (the height of the gray bars). (The vector \mathbf{p} also happens to be the projection of \mathbf{v} onto the ℓ_1 ball of radius $\eta\lambda$.)

Although this algorithm is less efficient than the quickselect-like algorithm when run in serial, the sort in line 2 and the cumulative sums in lines 4 and 6 (Ladner and Fischer, 1980) can be parallelized to run in $O(\log n)$ passes each.

4 Transformer

The Transformer network, introduced by Vaswani et al. (2017), is a sequence-to-sequence model in which both the encoder and the decoder consist of stacked self-attention layers. Each layer of the decoder can attend to the previous layer of the decoder and the output of the encoder. The multi-head attention uses two affine transformations, followed by a softmax. Additionally, each layer has a position-wise feed-forward neural network (FFN) with a hidden layer of rectified linear units:

$$\text{FFN}(x) = W_2(\max(0, W_1x + b_1)) + b_2.$$

The hidden layer size (number of columns of W_1) is typically four times the size of the model dimension. Both the multi-head attention and the feed-forward neural network have residual connections that allow information to bypass those layers.

4.1 Auto-sizing Transformer

Though the Transformer has demonstrated remarkable success on a variety of datasets, it is highly over-parameterized. For example, the English-German WMT '14 Transformer-base model proposed in Vaswani et al. (2017) has more




Figure 2: Architecture of the Transformer (Vaswani et al., 2017). We apply the auto-sizing method to the feed-forward (blue rectangles) and multi-head attention (orange rectangles) in all n layers of the encoder and decoder. Note that there are residual connections that can allow information and gradients to bypass any layer we are auto-sizing.

than 60M parameters. Whereas early NMT models such as Sutskever et al. (2014) have most of their parameters in the embedding layers, the added complexity of the Transformer, plus parallel developments reducing vocabulary size (Sennrich et al., 2016) and sharing embeddings (Press and Wolf, 2017) has shifted the balance. Nearly 31% of the English-German Transformer’s parameters are in the attention layers and 41% in the position-wise feed-forward layers.

Accordingly, we apply the auto-sizing method to the Transformer network, and in particular to the two largest components, the feed-forward layers and the multi-head attentions (blue and orange rectangles in Figure 2). A difference from the work of Murray and Chiang (2015) is that there are residual connections that allow information to bypass the layers we are auto-sizing. If the regularizer drives all the neurons in a layer to zero, information can still pass through. Thus, auto-sizing can effectively prune out an entire layer.

4.2 Random Search

As an alternative to grid-based searches, random hyperparameter search has been demonstrated to be a strong baseline for neural network architecture searches as it can search between grid points to increase the size of the search space (Bergstra

Dataset	Size
Ara-Eng	234k
Fra-Eng	235k
Hau-Eng	45k
Tir-Eng	15k

Table 1: Number of parallel sentences in training bi-texts. The French-English and Arabic-English data is from the 2017 IWSLT campaign (Mauro et al., 2012). The much smaller Hausa-English and Tigrinya-English data is from the LORELEI project.

and Bengio, 2012). In fact, Li and Talwalkar (2019) recently demonstrated that many architecture search methods do not beat a random baseline. In practice, randomly searching hyperparameter domains allows for an intuitive mixture of continuous and categorical hyperparameters with no constraints on differentiability (Maclaurin et al., 2015) or need to cast hyperparameter values into a single high-dimensional space to predict new values (Bergstra et al., 2011).

5 Experiments

All of our models are trained using the fairseq implementation of the Transformer (Gehring et al., 2017).² Our GPU-optimized, proximal gradient algorithms are implemented in PyTorch and are publicly available.³ For the random hyperparameter search experiments, we use SHADHO,⁴ which defines the hyperparameter tree, generates from it, and manages distributed resources (Kinnison et al., 2018). Our SHADHO driver file and modifications to fairseq are also publicly available.⁵

5.1 Settings

We looked at four different low-resource language pairs, running experiments in five directions: Arabic-English, English-Arabic, French-English, Hausa-English, and Tigrinya-English. The Arabic and French data comes from the IWSLT 2017 Evaluation Campaign (Mauro et al., 2012). The Hausa and Tigrinya data were provided by the LORELEI project with custom train/dev/test splits. For all languages, we tokenized and truecased the data using scripts from Moses (Koehn et al., 2007). For the Arabic systems, we translit-

erated the data using the Buckwalter transliteration scheme. All of our systems were run using subword units (BPE) with 16,000 merge operations on concatenated source and target training data (Sennrich and Haddow, 2016). We clip norms at 0.1, use label smoothed cross-entropy with value 0.1, and an early stopping criterion when the learning rate is smaller than 10^{-5} . All of our experiments were done using the Adam optimizer (Kingma and Ba, 2015), a learning rate of 10^{-4} , and dropout of 0.1. At test time, we decoded using a beam of 5 with length normalization (Boulanger-Lewandowski et al., 2013) and evaluate using case-sensitive, detokenized BLEU (Papineni et al., 2002).

5.1.1 Baseline

The originally proposed Transformer model is too large for our data size – the model will overfit the training data. Instead, we use the recommended settings in fairseq for IWSLT German-English as a baseline since two out of our four language pairs are also from IWSLT. This architecture has 6 layers in both the encoder and decoder, each with 4 attention heads. Our model dimension is $d_{model} = 512$, and our FFN dimension is 1024.

5.1.2 Auto-sizing parameters

Auto-sizing is implemented as two different types of group regularizers, $\ell_{2,1}$ and $\ell_{\infty,1}$. We apply the regularizers to the feed-forward network and multi-head attention in each layer of the encoder and decoder. We experiment across a range of regularization coefficient values, λ , that control how large the regularization proximal gradient step will be. We note that different regularization coefficient values are suited for different types or regularizers. Additionally, all of our experiments use the same batch size, which is also related to λ .

5.1.3 Random search parameters

As originally proposed, the Transformer network has 6 encoder layers, all identical, and 6 decoder layers, also identical. For our random search experiments, we sample the number of attention heads from $\{4, 8, 16\}$ and the model dimension (d_{model}) from $\{128, 256, 512, 1024, 2048\}$. Diverging from most implementations of the Transformer, we do not require the same number of encoder and decoder layers, but instead sample each from $\{2, 4, 6, 8\}$. Within a layer, we also sample the size of the feed-forward network (FFN), vary-

²<https://github.com/pytorch/fairseq>

³<https://github.com/KentonMurray/ProxGradPytorch>

⁴<https://github.com/jeffkinnison/shadho>

⁵https://bitbucket.org/KentonMurray/fairseq_autosizing

ing our samples over {512, 1024, 2048}. This too differs from most Transformer implementations, which have identical layer hyperparameters.

5.2 Auto-sizing vs. Random Search

Table 2 compares the performance of random search with auto-sizing across, BLEU scores, model size, and training times. The baseline system, the recommended IWSLT setting in fairseq, has almost 40 million parameters. Auto-sizing the feed-forward network sub-components in each layer of this baseline model with $\ell_{2,1} = 10.0$ or $\ell_{\infty,1} = 100.0$ removes almost one-third of the total parameters from the model. For Hausa-English and Tigrinya-English, this also results in substantial BLEU score gains, while only slightly hurting performance for French-English. The BLEU scores for random search beats the baseline for all language pairs, but auto-sizing still performs best on Tigrinya-English – even with 72 different, random hyperparameter configurations.

Auto-sizing trains in a similar amount of time to the baseline system, whereas the cumulative training time for all of the models in random search is substantially slower. Furthermore, for Tigrinya-English and French-English, random search found models that were almost 10 and 5 times larger respectively than the auto-sized models.

5.3 Training times

One of the biggest downsides of searching over architectures using a random search process is that it is very time and resource expensive. Contrary to that, auto-sizing relies on only training *one* model.

Auto-sizing relies on a proximal gradient step after a standard gradient descent step. However, the addition of these steps for our two group regularizers does not significantly impact training times. Table 3 shows the total training time for both $\ell_{2,1} = 0.1$ and $\ell_{\infty,1} = 0.5$. Even with the extra proximal step, auto-sizing using $\ell_{2,1}$ actually converges faster on two of the five language pairs. Note that these times are for smaller regularization coefficients. Larger coefficients will cause more values to go to zero, which will make the model converge faster.

5.4 Auto-sizing Sub-Components

As seen above, on very low-resource data, auto-sizing is able to quickly learn smaller, yet better, models than the recommended low-resource transformer architecture. Here, we look at the impact

of applying auto-sizing to various sub-components of the Transformer network. In section 3, following the work of Murray and Chiang (2015), auto-sizing is described as intelligently applying a group regularizer to our objective function. The relative weight, or regularization coefficient, is a hyperparameter defined as λ . In this section, we also look at the impact of varying the strength of this regularization coefficient.

Tables 4 and 5 demonstrate the impact of varying the regularization coefficient strength has on BLEU scores and model size across various model sub-components. Recall that each layer of the Transformer network has multi-head attention sub-components and a feed-forward network sub-component. We denote experiments only applying auto-sizing to feed-forward network as “FFN”. We also experiment with auto-sizing the multi-head attention in conjunction with the FFN, which we denote “All”. A regularization coefficient of 0.0 refers to the baseline model without any auto-sizing. Columns which contain percentages refer to the number of rows in a PyTorch parameter that auto-sizing was applied to, that were entirely driven to zero. In effect, neurons deleted from the model. Note that individual values in a row may be zero, but if even a single value remains, information can continue to flow through this and it is not counted as deleted. Furthermore, percentages refer only to the parameters that auto-sizing was applied to, not the entire model. As such, with the prevalence of residual connections, a value of 100% does not mean the entire model was deleted, but merely specific parameter matrices. More specific experimental conditions are described below.

5.4.1 FFN matrices and multi-head attention

Rows corresponding to “All” in tables 4 and 5 look at the impact of varying the strength of both the $\ell_{\infty,1}$ and $\ell_{2,1}$ regularizers across all learned parameters in the encoder and decoders (multi-head and feed-forward network parameters). Using $\ell_{\infty,1}$ regularization (table 5), auto-sizing beats the baseline BLEU scores on three language pairs: Hau–Eng, Tir–Eng, Fra–Eng. However, BLEU score improvements only occur on smaller regularization coefficients that do not delete model portions.

Looking at $\ell_{2,1}$ regularization across all learned parameters of both the encoder and decoder (“Enc+Dec All” in table 4), auto-sizing beats the baseline on four of the five language pairs (all except Eng–Ara). Again, BLEU gains are on smaller

Language Pair	Search Strategy	BLEU	Model Size	Training Time
Tir-Eng	Standard Baseline	3.6	39.6M	1.2k
	Random Search	6.7	240.4M	43.7k
	Auto-sizing $\ell_{\infty,1}$	7.5	27.1M	2.1k
	Auto-sizing $\ell_{2,1}$	7.4	27.1M	1.2k
Hau-Eng	Standard Baseline	13.9	39.6M	4.2k
	Random Search	17.2	15.4M	87.0k
	Auto-sizing $\ell_{\infty,1}$	15.0	27.1M	7.6k
	Auto-sizing $\ell_{2,1}$	14.8	27.1M	3.5k
Fra-Eng	Standard Baseline	35.0	39.6M	11.3k
	Random Search	35.3	121.2M	116.0k
	Auto-sizing $\ell_{\infty,1}$	34.3	27.1M	28.8k
	Auto-sizing $\ell_{2,1}$	33.5	27.1M	11.5k

Table 2: Comparison of BLEU scores, model size, and training time on Tigrinya-English, Hausa-English, and French-English. Model size is the total number of parameters. Training time is measured in seconds. Baseline is the recommended low-resource architecture in fairseq. Random search represents the best model found from 72 (Tigrinya), 40 (Hausa), and 10 (French) different randomly generated architecture hyperparameters. Both auto-sizing methods, on both languages, start with the exact same initialization and number of parameters as the baseline, but converge to much smaller models across all language pairs. On the very low-resource languages of Hausa and Tigrinya auto-sizing finds models with better BLEU scores. Random search is eventually able to find better models on French and Hausa, but is an order of magnitude slower.

Language Pair	Baseline	$\ell_{2,1}$	$\ell_{\infty,1}$
Fra-Eng	11.3k	11.5k	28.8k
Ara-Eng	15.1k	16.6k	40.8k
Eng-Ara	16.6k	11.0k	21.9k
Hau-Eng	4.2k	3.5k	7.6k
Tir-Eng	1.2k	1.2k	2.1k

Table 3: Overall training times in seconds on a Nvidia GeForce GTX 1080Ti GPU for small regularization values. Note that high regularization values will delete too many values and cause training to end sooner. In general, $\ell_{2,1}$ regularization does not appreciably slow down training, but $\ell_{\infty,1}$ can be twice as slow. Per epoch, roughly the same ratios in training times hold.

regularization coefficients, and stronger regularizers that delete parts of the model hurt translation quality. Multi-head attention is an integral portion of the Transformer model and auto-sizing this generally leads to performance hits.

5.4.2 FFN matrices

As the multi-head attention is a key part of the Transformer, we also looked at auto-sizing just the feed-forward sub-component in each layer of the encoder and decoder. Rows deonted by “FFN” in tables 4 and 5 look at applying auto-sizing to all

of the feed-forward network sub-components of the Transformer, but not to the multi-head attention. With $\ell_{\infty,1}$ regularization, we see BLEU improvements on four of the five language pairs. For both Hausa-English and Tigrinya-English, we see improvements even after deleting all of the feed-forward networks in all layers. Again, the residual connections allow information to flow around these sub-components. Using $\ell_{2,1}$ regularization, we see BLEU improvements on three of the language pairs. Hausa-English and Tigrinya-English maintain a BLEU gain even when deleting all of the feed-forward networks.

Auto-sizing only the feed-forward sub-component, and not the multi-head attention part, results in better BLEU scores, even when deleting all of the feed-forward network components. Impressively, this is with a model that has fully *one-third* fewer parameters in the encoder and decoder layers. This is beneficial for faster inference times and smaller disk space.

5.4.3 Encoder vs. Decoder

In table 4, experiments on Hau-Eng look at the impact of auto-sizing either the encoder or the decoder separately. Applying a strong enough regularizer to delete portions of the model ($\ell_{2,1} \geq 1.0$) only to the decoder (“Decoder All” and “Decoder

	Model Portion	$\ell_{2,1}$ coefficient									
		0.0	0.1	0.25	0.5	1.0	10.0				
Hau-Eng	Encoder All	13.9	16.0	17.1	17.4	15.3	89%	16.4	100%		
	Encoder FFN		15.4	15.1	16.3	15.9	100%	16.7	100%		
	Decoder All		12.6	16.1	16.2	13.0	3%	0.0	63%		
	Decoder FFN		11.8	14.7	14.4	11.7	79%	13.1	100%		
	Enc+Dec All		15.8	17.4	17.8	12.5	61%	0.0	100%		
	Enc+Dec FFN		14.7	15.3	14.2	12.8	86%	14.8	100%		
Tir-Eng	Encoder All	3.6	3.3	4.7	5.3	7.2		8.4	100%		
	Enc+Dec All		3.8	4.0	6.5	7.0		0.0	100%		
	Enc+Dec FFN		4.0	4.2	3.3	5.1		7.4	100%		
Fra-Eng	Encoder All	35.0	35.7	34.5	34.1	33.6	97%	32.8	100%		
	Enc+Dec All		35.2	33.1	29.8	23%	24.2	73%	0.3	100%	
	Enc+Dec FFN		35.6	35.0	34.2	15%	34.2	98%	33.5	100%	
Ara-Eng	Enc+Dec All	27.9	28.0	24.7	1%	20.9	20%	14.3	72%	0.3	100%
	Enc+Dec FFN		26.9	26.7	1%	25.5	23%	25.9	97%	25.7	100%
Eng-Ara	Enc+Dec All	9.4	8.7	7.5		5.8	23%	3.7	73%	0.0	100%
	Enc+Dec FFN		8.6	8.3	3%	8.3	22%	7.9	93%	8.0	100%

Table 4: BLEU scores and percentage of parameter rows deleted by auto-sizing on various sub-components of the model, across varying strengths of $\ell_{2,1}$ regularization. 0.0 refers to the baseline without any regularizer. Blank spaces mean less than 1% of parameters were deleted. In the two very low-resource language pairs (Hausa-English and Tigrinya-English), deleting large portions of the encoder can actually help performance. However, deleting the decoder hurts performance.

		$\ell_{\infty,1}$								
		0.0	0.1	0.25	0.5	1.0	10.0		100.0	
Hau-Eng	Enc+Dec All	13.9	15.5	14.7	16.0	16.7	14.9	4%	1.5	100%
	Enc+Dec FFN		13.4	14.3	14.1	12.9	15.3	0%	15.0	100%
Tir-Eng	Enc+Dec All	3.6	4.6	3.4	3.4	3.7	7.4	0%	2.4	100%
	Enc+Dec FFN		3.6	3.8	3.9	3.6	4.7	0%	7.5	100%
Fra-Eng	Enc+Dec All	35.0	35.2	35.4	34.9	35.3	26.3	13%	1.7	100%
	Enc+Dec FFN		34.8	35.5	35.4	35.0	34.1	0%	34.3	100%
Ara-Eng	Enc+Dec All	27.9	27.3	27.5	27.6	26.9	18.5	22%	0.6	100%
	Enc+Dec FFN		27.8	27.2	28.3	27.6	25.4	0%	25.4	100%
Eng-Ara	Enc+Dec All	9.4	9.1	8.3	8.4	8.7	5.2	25%	0.6	100%
	Enc+Dec FFN		8.8	9.2	9.0	8.9	8.2	0%	8.3	100%

Table 5: BLEU scores and percentage of model deleted using auto-sizing with various $\ell_{\infty,1}$ regularization strengths. On the very low-resource language pairs of Hausa-English and Tigrinya-English, auto-sizing the feed-forward networks of the encoder and decoder can improve BLEU scores.

FFN”) results in a BLEU score drop. However, applying auto-sizing to only the encoder (“Encoder All” and “Encoder FFN”) yields a BLEU gain while creating a smaller model. Intuitively, this makes sense as the decoder is closer to the output of the network and requires more modeling expressivity.

In addition to Hau–Eng, table 4 also contains experiments looking at auto-sizing all sub-components of all encoder layers of Tir–Eng and Fra–Eng. For all three language pairs, a small regularization coefficient for the $\ell_{2,1}$ regularizer applied to the encoder increases BLEU scores. However, no rows are driven to zero and the model size remains the same. Consistent with Hau–Eng, using a larger regularization coefficient drives all of the encoder’s weights to all zeros. For the smaller Hau–Eng and Tir–Eng datasets, this actually results in BLEU gains over the baseline system. Surprisingly, even on the Fra–Eng dataset, which has more than 15x as much data as Tir–Eng, the performance hit of deleting the *entire* encoder was only 2 BLEU points.

Recall from Figure 2 that there are residual connections that allow information and gradients to flow around both the multi-head attention and feed-forward portions of the model. Here, we have the case that all layers of the encoder have been completely deleted. However, the decoder still attends over the source word and positional embeddings due to the residual connections. We hypothesize that for these smaller datasets that there are too many parameters in the baseline model and over-fitting is an issue.

5.5 Random Search plus Auto-sizing

Above, we have demonstrated that auto-sizing is able to learn smaller models, faster than random search, often with higher BLEU scores. To compare whether the two architecture search algorithms (random and auto-sizing) can be used in conjunction, we also looked at applying both $\ell_{2,1}$ and $\ell_{\infty,1}$ regularization techniques to the FFN networks in all encoder and decoder layers *during* random search. In addition, this looks at how robust the auto-sizing method is to different initial conditions.

For a given set of hyperparameters generated by the random search process, we initialize three *identical* models and train a baseline as well as one with each regularizer ($\ell_{2,1} = 1.0$ and $\ell_{\infty,1} = 10.0$).

	none	$\ell_{2,1}$	$\ell_{\infty,1}$
Hau–Eng	17.2	16.6	17.8
Tir–Eng	6.7	7.9	7.6
Fra–Eng	35.4	34.7	34.1
Ara–Eng	27.6	25.6	25.9
Eng–Ara	9.0	7.6	8.4

Table 6: Test BLEU scores for the models with the best dev perplexity found using random search over number of layers and size of layers. Regularization values of $\ell_{2,1} = 1.0$ and $\ell_{\infty,1} = 10.0$ were chosen based on tables 4 and 5 as they encouraged neurons to be deleted. For the very low-resource language pairs, auto-sizing helped in conjunction with random search.

We trained 216 Tir–Eng models ($3 \cdot 72$ hyperparameter config.), 120 Hau–Eng, 45 Ara–Eng, 45 Eng–Ara, and 30 Fra–Eng models. Using the model with the best dev perplexity found during training, table 6 shows the test BLEU scores for each of the five language pairs. For the very low-resource language pairs of Hau–Eng and Tir–Eng, auto-sizing is able to find the best BLEU scores.

6 Conclusion

In this paper, we have demonstrated the effectiveness of auto-sizing on the Transformer network. On very low-resource datasets, auto-sizing was able to improve BLEU scores by up to 3.9 points while simultaneously deleting one-third of the parameters in the encoder and decoder layers. This was accomplished while being significantly faster than other search methods.

Additionally, we demonstrated how to apply proximal gradient methods efficiently using a GPU. Previous work on optimizing proximal gradient algorithms serious impacts speed performance when the computations are moved off of a CPU and parallelized. Leveraging sorting and prefix summation, we reformulated these methods to be GPU efficient.

Overall, this paper has demonstrated the efficacy of auto-sizing on a natural language processing application with orders of magnitude more parameters than previous work. With a focus on speedy architecture search and an emphasis on optimized GPU algorithms, auto-sizing is able to improve machine translation on very low-resource language pairs without being resource or time-consuming.

Acknowledgements

This research was supported in part by University of Southern California, subcontract 67108176 under DARPA contract HR0011-15-C-0115. We would like to thank Justin DeBenedetto for helpful comments.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proc. ICLR*.
- P. G. Benardos and G.-C. Vosniakos. 2007. Optimizing feedforward artificial neural network architecture. *Engineering Applications of Artificial Intelligence*, 20:365–382.
- James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13:281–305.
- James S. Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. 2011. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems*, pages 2546–2554.
- Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. 2013. Audio chord recognition with recurrent neural networks. In *Proc. International Society for Music Information Retrieval*, pages 335–340.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. 2010. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122.
- Denny Britz, Anna Goldie, Minh-Thang Luong, and Quoc Le. 2017. Massive exploration of neural machine translation architectures. In *Proc. EMNLP*, pages 1442–1451.
- Kai-Bo Duan and S. Sathya Keerthi. 2005. Which is the best multiclass SVM method? An empirical study. In *International Workshop on Multiple Classifier Systems*, pages 278–285.
- John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. 2008. Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In *Proc. ICML*, pages 272–279.
- Frauke Friedrichs and Christian Igel. 2005. Evolutionary tuning of multiple SVM parameters. *Neurocomputing*, 64:107–117.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional Sequence to Sequence Learning. In *Proc. ICML*.
- Diederik P. Kingma and Jimmy Lei Ba. 2015. Adam: A method for stochastic optimization. In *Proc. ICLR*.
- Jeffery Kinnison, Nathaniel Kremer-Herman, Douglas Thain, and Walter Scheirer. 2018. Shadho: Massively scalable hardware-aware distributed hyperparameter optimization. In *Proc. IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 738–747.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Proc. ACL: Demos*, pages 177–180.
- Richard E. Ladner and Michael J. Fischer. 1980. [Parallel prefix computation](#). *J. ACM*, 27(4):831–838.
- Liam Li and Ameet Talwalkar. 2019. Random search and reproducibility for neural architecture search.
- Dougal Maclaurin, David Duvenaud, and Ryan Adams. 2015. Gradient-based hyperparameter optimization through reversible learning. In *Proc. ICML*, pages 2113–2122.
- Cettolo Mauro, Girardi Christian, and Federico Marcello. 2012. Wit3: Web inventory of transcribed and translated talks. In *Proc. EAMT*, pages 261–268.
- Kenton Murray and David Chiang. 2015. Auto-sizing neural networks: With applications to n -gram language models. In *Proc. EMNLP*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proc. ACL*, pages 311–318.
- Neal Parikh and Stephen Boyd. 2014. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):123–231.
- Martin Popel and Ondřej Bojar. 2018. Training tips for the Transformer model. *The Prague Bulletin of Mathematical Linguistics*, 110(1):43–70.
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models](#). In *Proc. EACL: Volume 2, Short Papers*, pages 157–163.
- Ariadna Quattoni, Xavier Carreras, Michael Collins, and Trevor Darrell. 2009. An efficient projection for $l_{1,\infty}$ regularization. In *Proc. ICML*, pages 857–864.
- Rico Sennrich and Barry Haddow. 2016. Linguistic input features improve neural machine translation. In *Proc. First Conference on Machine Translation: Volume 1, Research Papers*, volume 1, pages 83–91.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proc. ACL*, pages 1715–1725.

- Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. 2012. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, pages 2951–2959.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and policy considerations for deep learning in NLP. In *Proc. ACL*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*, pages 3104–3112.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Aaron Vose, Jacob Balma, Alex Heye, Alessandro Rigazzi, Charles Siegel, Diana Moise, Benjamin Robbins, and Rangan Sukumar. 2019. Recombination of artificial neural networks. arXiv:1901.03900.